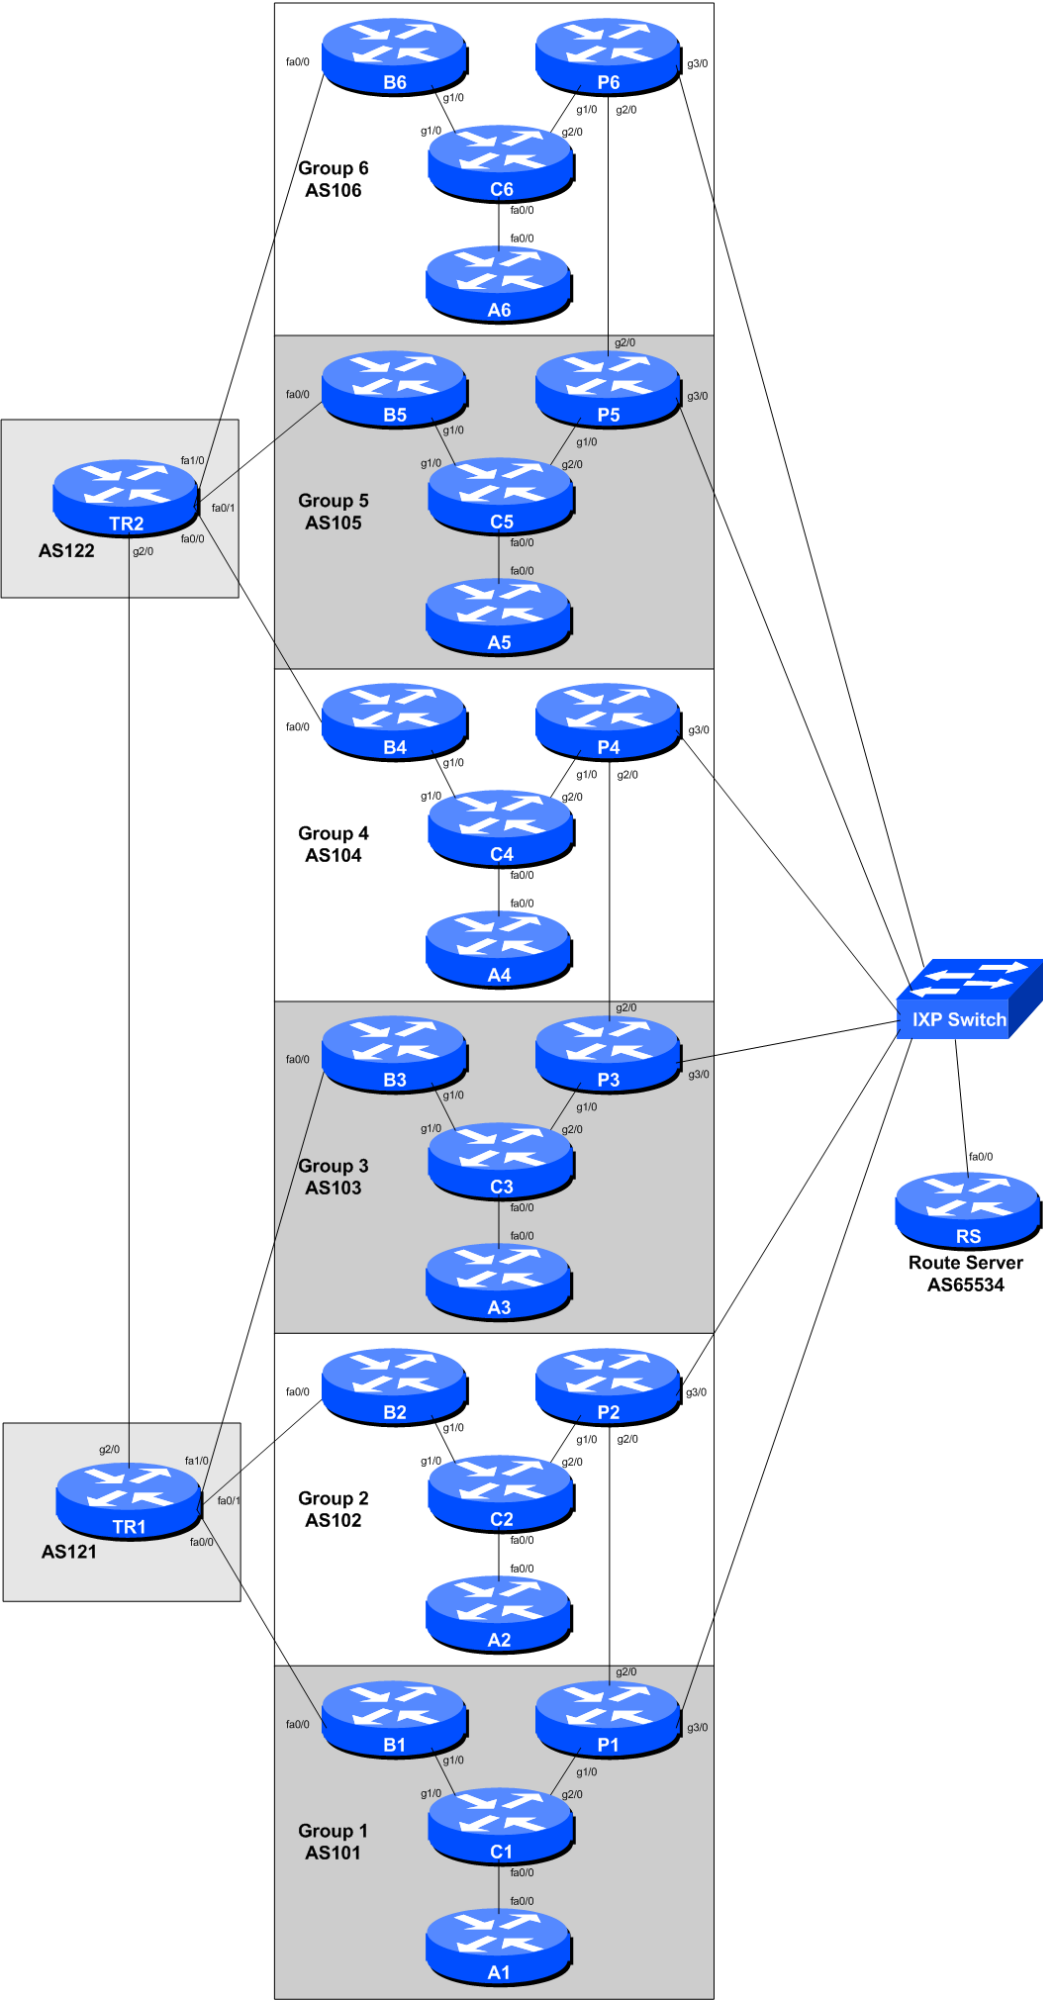


Peering & IXP Lab (Part 3: IXP)

The purpose of this part of the lab is introduce an Internet Exchange Point into our lab. IXPs are a very important if not critical component of today's Internet architecture, and it is vitally important to ensure the correct configuration so that network operators gain maximum advantage from their participation at an IXP.

Lab Topology

The lab topology has been further enhanced according to the diagram below. An Internet Exchange Point and its router server are now installed.



Configuring the IXP links

Each group should now configure their link to the IXP according to the above diagram.

Physical Link

Consult the [Address Plan](#) document for the address space used by the IXP. Following the document, configure the interface on the router accordingly.

```
interface GigabitEthernet 3/0
  description Link to IXP
  ip address 100.127.1.5 255.255.255.0
  ipv6 address 2001:DB8:FFFF:1::5/64
  !
```

Note the subnet masks - this time the ethernet is **NOT** a point-to-point link but a shared LAN media. Once the interfaces have been configured see if you can ping any of the other groups on their IXP addresses (both IPv4 and IPv6). Are you able to ping the Route Server too?

Configuring IS-IS

Do not configure IS-IS towards any IXP peer! They are not part of your autonomous system.

However, so that traceroutes across the IXP do not break, we might wish to carry the IXP LAN address block within our IS-IS (**not** iBGP). To do this, we simply mark the IXP facing interface as passive in the IS-IS configuration. Here is an example for AS106:

```
router isis as106
  passive-interface GigabitEthernet 3/0
```

If you recall from the IS-IS presentation, this will tell IS-IS to announce the subnet attached to this interface.

Now all routers in your AS will see the IXP LAN address - check from your Core and Border routers, just to make sure.

Configuring eBGP

We now configure eBGP with the Exchange Point's Route Server (we might add in bi-lateral BGP peering later, but for now we will just peer with the Route Server).

The Route Server sits in AS 65534 - this is a private AS, and is not visible on the public Internet. In fact, we don't want this AS to be visible inside our own AS either, and that's one of the unique features of a Route Server - it does not add its AS into the AS path when distributing prefixes to its eBGP neighbours.

DO NOT forget to filter what you hear from the Route Server, and what you send to the Route Server. You should only accept their address blocks from the other IXP participants (they may send you more by mistake!), and you should only send prefixes you originated!

Here is a configuration example for AS104 - note that we are reusing some configuration we have set up earlier:

```
ip prefix-list AS104-block permit 100.68.40.0/24
ipv6 prefix-list AS104-v6block permit 2001:DB8:40::/48
!
ip prefix-list IXP-RS permit 100.68.10.0/24
ip prefix-list IXP-RS permit 100.68.20.0/24
ip prefix-list IXP-RS permit 100.68.30.0/24
ip prefix-list IXP-RS permit 100.68.40.0/24
ip prefix-list IXP-RS permit 100.68.50.0/24
ip prefix-list IXP-RS permit 100.68.60.0/24
!
ipv6 prefix-list IXP-v6RS permit 2001:DB8:10::/48
ipv6 prefix-list IXP-v6RS permit 2001:DB8:20::/48
ipv6 prefix-list IXP-v6RS permit 2001:DB8:30::/48
ipv6 prefix-list IXP-v6RS permit 2001:DB8:40::/48
ipv6 prefix-list IXP-v6RS permit 2001:DB8:50::/48
ipv6 prefix-list IXP-v6RS permit 2001:DB8:60::/48
!
router bgp 104
  address-family ipv4
    neighbor 100.127.1.254 remote-as 65534
    neighbor 100.127.1.254 description eBGP with IXP RS
    neighbor 100.127.1.254 password ixp-rs
    neighbor 100.127.1.254 prefix-list AS104-block out
    neighbor 100.127.1.254 prefix-list IXP-RS in
  !
  address-family ipv6
    neighbor 2001:DB8:FFFF:1::FE remote-as 65534
    neighbor 2001:DB8:FFFF:1::FE description eBGP with IXP RS
    neighbor 2001:DB8:FFFF:1::FE password ixp-rs
    neighbor 2001:DB8:FFFF:1::FE prefix-list AS104-v6block out
    neighbor 2001:DB8:FFFF:1::FE prefix-list IXP-v6RS in
  !
```

Once this has been configured, has the BGP session with the Route Server established? If not, why not? What do the router logs tell you?

You will notice from the logs that the router is complaining about a BGP peer AS not being in the announced AS path - this is Cisco IOS protecting against improper BGP announcements, as according to the BGP RFC, the AS PATH of the neighbouring AS must appear as the adjacent AS in the AS PATH. And if you recall from early in the notes, that was a special feature of the Route Server: its AS does not appear in the path.

So we need to turn this safety check off in IOS:

```
router bgp 105
no bgp enforce-first-as
```

and once this has been done you will now see that the eBGP session with the Route Server will have been established.

What do you now see in the BGP table?

What about the routes between you and your private peer that you set up earlier? Which is the best path now? Through the IXP, or over the private peering link?

Private Peering link versus IXP Peering

We are now going to deal with the issue where we see two paths between us and our private peer. One is via our private peering link, the other is via our peering with them across the IXP.

In day to day Internet operations, network operators prioritise links according to the value they attach to them - the list goes something like this:

Type of Link	Priority (local preference)
Private Peer	200
IXP Bi-lateral	180
IXP RS	170
Customer	120
(default)	100
Local Transit	70
Global Transit	50

Obviously there will be many variations on this theme, but the principle remains the same. Peering links have no operational cost, so are highly preferred over links which have an operational cost (transit). Private peering links are preferred over IXP links as the former is brokered directly with the partner, while the IXP links are via a third party infrastructure. It is not physically possible to peer privately with every operator, and this is the function that the IXP then provides (as was covered in the course presentations).

We will now attach local preference to the routes we hear from our private peer and from the IXP Route Server, according to the table above.

To do this we will create two route-maps, one for the private peer, the other for the RS peering. Here is an example for AS101:

```
route-map private-peer-in permit 10
set local-preference 200
!
route-map IXP-RS-peer-in permit 10
set local-preference 170
!
router bgp 101
address-family ipv4
```

```
neighbor 100.68.10.26 route-map private-peer-in in
neighbor 100.127.1.254 route-map IXP-RS-peer-in in
!
address-family ipv6
neighbor 2001:DB8:10:12::1 route-map private-peer-in in
neighbor 2001:DB8:FFFF:1::FE route-map IXP-RS-peer-in in
!
```

Once this is configured, do a route-refresh inbound on the two peering, and now you should now see the local preferences attached to the routes from the IXP and from the private peer. What has happened now?

Confirmation

Check on the Border, Core, Access and Peering Routers for what you now see in the BGP table.

What is the best path to your private peer? What does trace route tell you?

Hopefully you will see that the best path to your private peer will be via the private peering link. And the routes to the rest of the class will be via the Internet Exchange Point. The only traffic going via the Upstream Provider now will be traffic out to the Internet itself. If this is not the case, you will need to start doing some troubleshooting!

Peering with Route Server or Peering Directly with IXP peers?

The final part of this workshop lab is to investigate how to set up peering directly with IXP Peers. There are three types of peering policies adopted by network operators today:

Peering Policy	Description
<i>Open</i>	Network Operator will peer with allcomers, no questions asked. At an IXP this means they will peer with the Route Server.
<i>Selective</i>	Network Operator will usually peer with most operators, but enters a conversation with the peering partner first before establishing the link. At an IXP this means they will set up a direct peering across the IX fabric.
<i>Restricted</i>	Network Operator will choose who they peer with, under very stringent conditions. They rarely show up at an IXP, and if they do, peering will be directly across the IX fabric.

We have set up our peering at the moment to assume that all groups have an *Open* peering policy. But what if they had a *Selective* policy instead? How do we configure that?

Bi-lateral peering across IX Fabric

What we will do now is modify our eBGP at the IXP so that we also include a direct eBGP session with our IXP peers. We'll set this up to supplement the Route Server (or we could simply remove the Route

Server peering once we have peered with all members of the IXP).

Here is a configuration example for AS102, peering with AS103 (again noting that we are re-using configuration created earlier on):

```
ip prefix-list AS102-block permit 100.68.20.0/24
ipv6 prefix-list AS102-v6block permit 2001:DB8:20::/48
!
ip prefix-list AS103-block permit 100.68.30.0/24
ipv6 prefix-list AS103-v6block permit 2001:DB8:30::/48
!
route-map IXP-bilateral-in permit 10
  set local-preference 170
!
router bgp 102
  address-family ipv4
    neighbor 100.127.1.3 remote-as 103
    neighbor 100.127.1.3 description eBGP with AS103
    neighbor 100.127.1.3 password cisco
    neighbor 100.127.1.3 prefix-list AS102-block out
    neighbor 100.127.1.3 prefix-list AS103-block in
    neighbor 100.127.1.3 route-map IXP-bilateral-in in
  !
  address-family ipv6
    neighbor 2001:DB8:FFFF:1::3 remote-as 103
    neighbor 2001:DB8:FFFF:1::3 description eBGP with AS103
    neighbor 2001:DB8:FFFF:1::3 password cisco
    neighbor 2001:DB8:FFFF:1::3 prefix-list AS102-v6block out
    neighbor 2001:DB8:FFFF:1::3 prefix-list AS103-v6block in
    neighbor 2001:DB8:FFFF:1::3 route-map IXP-bilateral-in in
  !
```

What do you see now?

You should see two paths to your IXP peers - they are almost indistinguishable apart from the router-id of the neighbouring router - one will be of the route-server, the other will be of the direct eBGP peer.

Repeat the above for all the members of the IXP.

Once you are peering with all of the members of the IXP, you can remove your peering with the Route Server if you wish:

- if you have decided that your group has a Selective Peering Policy, you may not want to peer with the Route Server, as then you won't be able to select the new peers (unless the Route Server operator offers a more sophisticated configuration than we have here in the lab)
- Many operators bi-lateral peer across the IXP but retain the Route Server peering as back as well - for redundancy purposes.

Appendix - Route Server Configuration

This appendix shows the configuration of the route server used for this workshop. It is Cisco IOS based - most route servers today run either on BIRD or a modified version of Quagga.

```
interface FastEthernet0/0
  description IXP LAN
  ip address 100.127.1.254 255.255.255.0
  ipv6 address 2001:DB8:FFFF:1::FE/64
!
router bgp 65534
  bgp log-neighbor-changes
  bgp deterministic-med
  no bgp default ipv4-unicast
  neighbor ixp-peers peer-group
  neighbor ixp-peers password ixp-rs
  neighbor v6ixp-peers peer-group
  neighbor v6ixp-peers password ixp-rs
  neighbor 100.127.1.1 remote-as 101
  neighbor 100.127.1.1 peer-group ixp-peers
  neighbor 100.127.1.1 description AS101 peer
  neighbor 100.127.1.2 remote-as 102
  neighbor 100.127.1.2 peer-group ixp-peers
  neighbor 100.127.1.2 description AS102 peer
  neighbor 100.127.1.3 remote-as 103
  neighbor 100.127.1.3 peer-group ixp-peers
  neighbor 100.127.1.3 description AS103 peer
  neighbor 100.127.1.4 remote-as 104
  neighbor 100.127.1.4 peer-group ixp-peers
  neighbor 100.127.1.4 description AS104 peer
  neighbor 100.127.1.5 remote-as 105
  neighbor 100.127.1.5 peer-group ixp-peers
  neighbor 100.127.1.5 description AS105 peer
  neighbor 100.127.1.6 remote-as 106
  neighbor 100.127.1.6 peer-group ixp-peers
  neighbor 100.127.1.6 description AS106 peer
  neighbor 2001:DB8:FFFF:1::1 remote-as 101
  neighbor 2001:DB8:FFFF:1::1 peer-group v6ixp-peers
  neighbor 2001:DB8:FFFF:1::1 description AS101 peer
  neighbor 2001:DB8:FFFF:1::2 remote-as 102
  neighbor 2001:DB8:FFFF:1::2 peer-group v6ixp-peers
  neighbor 2001:DB8:FFFF:1::2 description AS102 peer
  neighbor 2001:DB8:FFFF:1::3 remote-as 103
  neighbor 2001:DB8:FFFF:1::3 peer-group v6ixp-peers
  neighbor 2001:DB8:FFFF:1::3 description AS103 peer
  neighbor 2001:DB8:FFFF:1::4 remote-as 104
  neighbor 2001:DB8:FFFF:1::4 peer-group v6ixp-peers
  neighbor 2001:DB8:FFFF:1::4 description AS104 peer
  neighbor 2001:DB8:FFFF:1::5 remote-as 105
```

```
neighbor 2001:DB8:FFFF:1::5 peer-group v6ixp-peers
neighbor 2001:DB8:FFFF:1::5 description AS105 peer
neighbor 2001:DB8:FFFF:1::6 remote-as 106
neighbor 2001:DB8:FFFF:1::6 peer-group v6ixp-peers
neighbor 2001:DB8:FFFF:1::6 description AS106 peer
!
address-family ipv4
  neighbor ixp-peers route-server-client
  neighbor 100.127.1.1 activate
  neighbor 100.127.1.2 activate
  neighbor 100.127.1.3 activate
  neighbor 100.127.1.4 activate
  neighbor 100.127.1.5 activate
  neighbor 100.127.1.6 activate
  distance bgp 200 200 200
exit-address-family
!
address-family ipv6
  neighbor v6ixp-peers route-server-client
  neighbor 2001:DB8:FFFF:1::1 activate
  neighbor 2001:DB8:FFFF:1::2 activate
  neighbor 2001:DB8:FFFF:1::3 activate
  neighbor 2001:DB8:FFFF:1::4 activate
  neighbor 2001:DB8:FFFF:1::5 activate
  neighbor 2001:DB8:FFFF:1::6 activate
  distance bgp 200 200 200
exit-address-family
!
ip route 0.0.0.0 0.0.0.0 Null0
ipv6 route ::/0 Null0
```

From:

<https://workshops.nsrc.org/dokuwiki/> - **Workshops**

Permanent link:

<https://workshops.nsrc.org/dokuwiki/2016/pacnog19-routing/peering-ixp-part3>

Last update: **2016/11/30 23:02**

